Opponent Modelling in Automated Multi-Issue Negotiation Using Bayesian Learning[1]

Koen Hindriks Dmytro Tykhonov

Man-Machine Interaction Group,
Delft University of Technology,
Mekelweg 4, 2628 CD Delft

1 Introduction

In bilateral negotiation, two parties aim at reaching a joint agreement. They do so by exchanging various offers or bids using e.g. an alternating offers protocol [2]. In reaching such an agreement both parties usually aim to satisfy their own interests as best as possible, but have to take their opponent's preferences into account as well to reach an agreement at all. This is complicated by the fact that negotiating parties are generally not willing to reveal their preferences in order to avoid exploitation. As a result, both parties have incomplete information which makes it hard to decide on a good negotiation move and hard to reach an optimal agreement.

In this paper, we show that it is nonetheless possible to construct an *opponent model*, i.e. a model of the opponent's preferences that can be effectively used to improve negotiation outcomes. We provide a generic framework for learning both the preferences associated with issue values as well as the weights that rank the importance of issues to an agent. The main idea is to exploit certain structural features and rationality principles to guide the learning process and focuses the algorithm on the most likely preference profiles of an opponent. We present a learning algorithm based on Bayesian learning techniques that computes the probability that an opponent has a particular preference profile. Our approach can be integrated into various negotiating agents using different strategies.

2 Learning an Opponent Model

Our goal is to introduce a learning approach that can be used to model an opponent in a negotiation with imperfect information. First, we define a hypothesis space that specifies the range of opponent profiles that are considered. We do so by introducing various reasonable assumptions about the structure of opponent profiles as well as about an opponent's negotiation strategy.

Our first assumption is a common one, see e.g., [2], and assumes that the utility of a bid can be computed as a weighted sum of the utilities associated with the values for each issue. Utility functions modelling the preferences of an agent thus are linearly additive functions and are defined by a set of weights (or priorities) and corresponding evaluation functions for each of n issues. In order to learn an opponent's preference profile or utility function we need to learn both the issue weights as well as the evaluation functions. The objective of learning an opponent model thus is to find a model that is the most plausible candidate or best approximation of the opponent's preference profile.

Our next assumption concerns the issue priorities in a preference profile. We define the set of hypotheses about the private weights of an opponent as the set of all possible rankings of weights. It is then straightforward to associate real-valued numbers again with a hypothesis about weights, which can be computed as a linear function of the rank. Finally, we need to impose some additional structure on the evaluation functions

[1] The full version of this paper appeared in: *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16, 2008, Estoril, Portugal.*

in order to be able to learn a preference profile. To facilitate the learning of an opponent's preferences over issue values we introduce a hypothesis space of predefined function shapes (*downhill, uphill, and triangular*). To summarize, the set of hypotheses concerning an opponent's preference profile is a Cartesian product of the hypotheses about issue weights and shapes of issue evaluation functions.

The idea is to learn an opponent preference profile from its negotiation moves, i.e. the bids it proposes during a negotiation. In a Bayesian learning approach, this means we need to be able to update the probability associated with all hypotheses given new evidence, i.e. one of the bids. More precisely, we need to compute the probability of every hypothesis given the bid proposed by the opponent. In order to be able to use Bayes' rule to do this, however, we need some information about the utility that the opponent associates with a bid. To this end it is assumed that an agent's tactics during a negotiation can be defined by a monotonically decreasing function [4]. This assumption is reasonable given that negotiating agents have to concede to reach an agreement and still allows the use of various kinds of tactics since no exact knowledge about an opponent's negotiation tactics is assumed. More specifically, the rationality assumption is modelled as a probability distribution associated with a range of tactics. As a result, each utility associated with an opponent's bid thus also has an associated probability. Finally, during a negotiation an agent can use the updated probability distribution to compute expected utility of counteroffers it considers and choose one that e.g. maximizes the utility of its opponent, to increase the likelihood of acceptance by that opponent.

Experiments have been performed to show the effectiveness of our approach to learn the opponent model and to use it to find a good counteroffer (see the full paper). The Bayesian learning agents used in the experiment update their opponent model each time a new bid is received from the opponent in line with the Bayesian learning approach introduced above. The strategy used by the Bayesian learning agents is based on the smart meta-strategy of [1]. The agent starts with proposing a bid that has maximal utility given its own preferences. Each consecutive turn the agent can either accept the opponent's bid or send a counter-offer. The agent accepts a bid from its opponent when the utility of that bid is higher than the utility of its own last bid or the utility of the bid it would otherwise propose next. Otherwise, the agent will propose a counter-offer. In this domain, the Bayesian agents very efficiently learn issue weights when they are provided with domain knowledge, indicated by the fact that the negotiation trace almost coincides with the Pareto frontier. But even without domain knowledge the Bayesian agent needs little time to learn the issue evaluation functions and consecutively improves the weight estimations. The influence of the negotiation domain, preference profile, and opponent's strategy on the quality of learning was investigated in [3].

3 Conclusions

In this paper, an opponent modelling framework for bilateral multi-issue negotiation has been presented. The main idea proposed here to make opponent modelling in negotiation feasible is to assume that certain structural requirements on preference profiles and on the strategy of an opponent are in place. Due to the probabilistic nature of the model, these assumptions still allow for a great diversity of potential opponent models. The learning approach has been tested on several domains to demonstrate the effectiveness of the approach. The results moreover showed the effectiveness of using an opponent model in a negotiation strategy to improve the efficiency of the bidding process. In future work we will analyze the quality of the learned opponent model with respect to the original preferences profile of the opponent.

The learning approach does not rely on prior knowledge about e.g. the domain, but if such knowledge is available it can be incorporated and used to initialize probability distributions in the opponent model. Domain knowledge would also be useful to increase the efficiency of learning a correct opponent model in learning algorithm proposed (for more details see the full paper).

References

[1] Faratin, P., Sierra, C., and Jennings, N. R. 2003. Using Similarity Criteria to Make Negotiation Trade-Offs, Journal of Artificial Intelligence, 142 (2), pp. 205-237.

[2] Raiffa, H. *The Art and Science of Negotiation*, Harvard University Press. 1982.

[3] Hindriks, K.V., Tykhonov, D. 2008. Towards a Quality Assessment Method for Learning Preference Profiles in Negotiation, in Proc. of AMEC'08.

[4] Zeng, D., and Sycara, K. 1998. Bayesian Learning in Negotiation, International Journal of Human Computer Systems, vol. 48, pp. 125-141.